

IN THE CLAIMS:

- 1 1. (Previously Presented) A system for replay of a backup memory in a storage system
2 having a file system for managing transfer of data to and from an attached disk array, the
3 system comprising:
4 a log in the backup memory containing the storage system transaction entries ac-
5 cumulated after a consistency point at which time results of the storage system transac-
6 tion entries are committed to the disk array;
7 an initiator process that establishes a swarm of messages with respect to the stor-
8 age system transaction entries and delivers the swarm to the file system; and
9 a parallel disk information-retrieval process in the file system that is carried out in
10 an arbitrary order on the swarm of messages in parallel by a processor within the storage
11 system, wherein the processor commingles steps of each message to process the mes-
12 sages in parallel.
- 1 2. (Original) The system as set forth in claim 1 wherein each of the messages of the
2 swarm is identified by a transaction block including a pointer to one of the transaction
3 request entries in the log, respectively, and a state that indicates whether each of the mes-
4 sages is one of (a) newly transferred to the file system, (b) subject to completion of a
5 LOAD phase thereon by the disk information-retrieval process, (c) subject to completion
6 of a MODIFY phase thereon by a MODIFY process of the file system or (d) incapable of
7 being subject to the LOAD phase until a prerequisite event occurs.
- 1 3. (Original) The system as set forth in claim 2 wherein the prerequisite event is com-
2 pletion of the LOAD phase and a MODIFY phase with respect to another of the mes-
3 sages.

1 4. (Original) The system as set forth in claim 3 wherein the initiator process is adapted
2 to retransfer each of the messages incapable of being subject to a load phase until the pre-
3 requisite event occurs to the file system for completion of the LOAD phase after the pre-
4 requisite event occurs, respectively.

1 5. (Original) The system as set forth in claim 4 wherein the initiator is adapted to estab-
2 lish a skip state with respect to skipped messages for which a portion of the disk array
3 associated therewith is unavailable, the skip state thereby omitting the skipped messages
4 from the swarm.

1 6. (Original) The system as set forth in claim 4 wherein the file system includes a panic
2 state adapted to alert an operator if a first message received from the initiator in the
3 swarm is a message incapable of being subject to a load phase until a prerequisite event
4 occurs.

1 7. (Original) The system as set forth in claim 4 wherein the file system includes a panic
2 state adapted to alert an operator if a message retransferred by the initiator process is a
3 message incapable of being subject to a load phase until a prerequisite event occurs.

1 8. (Original) The system as set forth in claim 1 wherein the backup memory comprises
2 a non-volatile random access memory (NVRAM).

1 9. (Original) The system as set forth in claim 1 wherein the storage system comprises a
2 network storage appliance.

1 10. (Previously Presented) A method for replay of a backup memory in a storage system
2 having a file system for managing transfer of data to and from an attached disk array, the
3 method comprising:

4 accumulating, in a log in the backup memory, storage system transaction request
5 entries after a consistency point at which time results of the transaction request entries are
6 committed to the disk array;

7 establishing a swarm of messages with respect to the transaction request entries
8 and delivering the swarm to the file system; and

9 performing a parallel disk information-retrieval process of the file system on the
10 swarm of messages in parallel by one or more processors within the storage system.

1 11. (Original) The method as set forth in claim 10 further comprising establishing, for
2 each of the messages of the swarm, a transaction block including a pointer to one of the
3 transaction request entries in the log, respectively, and a state that indicates whether each
4 of the messages is one of (a) newly transferred to the file system, (b) subject to comple-
5 tion of a LOAD phase thereon by the disk information-retrieval process, (c) subject to
6 completion of a MODIFY phase thereon by a MODIFY process of the file system or (d)
7 incapable of being subject to the LOAD phase until a prerequisite event occurs.

1 12. (Original) The method as set forth in claim 11 wherein the prerequisite event is com-
2 pletion of the LOAD phase and a MODIFY phase with respect to another of the mes-
3 sages.

1 13. (Original) The method as set forth in claim 12 further comprising retransferring each
2 of the messages incapable of being subject to a load phase until the prerequisite event oc-
3 curs to the file system for completion of the LOAD phase after the prerequisite event oc-
4 curs, respectively.

1 14. (Original) The method as set forth in claim 10 wherein the storage system comprises
2 a network storage appliance.

1 15. (Previously Presented) A computer-readable medium including program instructions
2 executing on a computer for parallelized replay of a backup memory in a storage system
3 having a file system for managing transfer of data to and from an attached disk array, the
4 program instructions performing the steps of:

5 accumulating, in a log in the backup memory, storage system transaction request
6 entries after a consistency point at which results of the transaction request entries are
7 committed to the disk array;

8 establishing a swarm of messages with respect to the transaction request entries
9 and delivering the swarm to the file system; and

10 performing parallel a disk information-retrieval process of the file system on the
11 swarm of messages in parallel by one or more processors within the storage system.
12

1 16. (Original) The computer-readable medium as set forth in claim 15 further comprising
2 establishing, for each of the messages of the swarm, a transaction block including a
3 pointer to one of the transaction request entries in the log, respectively, in the log and a
4 state that indicates whether each of the messages is one of (a) newly transferred to the file
5 system, (b) subject to completion of the LOAD phase thereon by the disk information-
6 retrieval process, (c) subject to completion of a MODIFY phase thereon by a MODIFY
7 process of the file system or (d) incapable of being subject to the LOAD phase until a
8 prerequisite event occurs.

1 17. (Original) The computer-readable medium as set forth in claim 16 wherein the pre-
2 requisite event is completion of the LOAD phase and a MODIFY phase with respect to
3 another of the messages.

1 18. (Original) The computer-readable medium as set forth in claim 17 further comprising
2 retransferring each of the messages incapable of being subject to a load phase until the

3 prerequisite event occurs to the file system for completion of the LOAD phase after the
4 prerequisite event occurs, respectively.

1 19. (Original) The computer-readable medium as set forth in claim 15 wherein the stor-
2 age system comprises a network storage appliance.

1 20. (Previously Presented) An apparatus for replay of a backup memory in a storage sys-
2 tem having a file system for managing transfer of data to and from an attached disk array,
3 comprising:

4 a processor to determine a consistency point in time, said apparatus containing at
5 least one transaction entry accumulated after the consistency point, where at the time of
6 the consistency point the transaction entries are committed to the disk array;

7 a plurality of messages, each message of said plurality of messages being related
8 to a transaction entry of said transaction entries accumulated after the consistency point,
9 said plurality of messages being referred to as a swarm of messages;

10 an initiator process to deliver the swarm of messages to the file system; and

11 a parallel disk information-retrieval process that processes the swarm of messages
12 in parallel by one or more processors within the storage system.

1 21. (Previously Presented) The apparatus as set forth in claim 20, further comprising:

2 each of the messages of the swarm is identified by a transaction block including a
3 pointer to one of the transaction request entries.

1 22. (Previously Presented) The apparatus as set for in claim 20, further comprising:

2 a state that indicates whether each of the messages is one of

3 (a) newly transferred to the file system,

4 (b) subject to completion of a LOAD phase thereon by the disk informa-
5 tion-retrieval process,

6 (c) subject to completion of a MODIFY phase thereon by a MODIFY pro-
7 cess of the file system, or

8 (d) incapable of being subject to the LOAD phase until a prerequisite
9 event occurs.

1 23. (Previously Presented) The apparatus as set forth in claim 22, further comprising:
2 the prerequisite event is completion of the LOAD phase and a MODIFY phase
3 with respect to another of the messages.

1 24. (Previously Presented) The apparatus as set forth in claim 23, further comprising:
2 the initiator process is adapted to retransfer each of the messages incapable of be-
3 ing subject to a load phase until the prerequisite event occurs to the file system for com-
4 pletion of the LOAD phase after the prerequisite event occurs, respectively.

1 25. (Previously Presented) The apparatus as set forth in claim 20, further comprising:
2 the initiator is adapted to establish a skip state with respect to skipped messages
3 for which a portion of the disk array associated therewith is unavailable, the skip state
4 thereby omitting the skipped messages from the swarm.

1 26. (Previously Presented) A method for replay of a backup memory in a storage system
2 having a file system for managing transfer of data to and from an attached disk array, the
3 method comprising:
4 accumulating one or more transaction request entries after a consistency point,
5 said consistency point is a time at which results of the transaction request entries are
6 committed to the disk array;
7 establishing a plurality of messages with respect to the transaction request entries,
8 said plurality of messages being referred to as a swarm of messages and delivering the
9 swarm to the file system; and

10 executing a parallel disk information-retrieval process on the swarm of messages
11 in parallel by one or more processors within the storage system.

1 27. (Previously Presented) The method as set forth in claim 26, further comprising:
2 establishing, for each of the messages of the swarm, a transaction block including
3 a pointer to one of the transaction request entries in the log.

1 28. (Previously Presented) The method as set forth in claim 20, further comprising:
2 establishing a state that indicates whether each of the messages is one of
3 (a) newly transferred to the file system,
4 (b) subject to completion of a LOAD phase thereon by the disk informa-
5 tion-retrieval process,
6 (c) subject to completion of a MODIFY phase thereon by a MODIFY pro-
7 cess of the file system, or
8 (d) incapable of being subject to the LOAD phase until a prerequisite
9 event occurs.

1 29. (Previously Presented) The method as set forth in claim 28, further comprising:
2 using as the prerequisite event completion of the LOAD phase and a MODIFY
3 phase with respect to another of the messages.

1 30. (Previously Presented) The method as set forth in claim 29, further comprising:
2 retransferring each of the messages incapable of being subject to a load phase un-
3 til the prerequisite event occurs to the file system for completion of the LOAD phase af-
4 ter the prerequisite event occurs.

1 31. -35. (Cancelled)

1 36. (Previously Presented) A computer readable media, comprising:
2 said computer readable media having instructions written thereon for execution on
3 a processor for the practice of the method of claim 10 or claim 26.

1 37. (Cancelled)

1 38. (Previously Presented) The system of claim 1, further comprising:
2 a third process that modifies at least some messages in the swarm of messages
3 based on the order in which storage system transaction entries were stored in the log.

1 39. (Previously Presented) The method of claim 10, further comprising:
2 modifying at least some messages in the swarm of messages based on the order in
3 which storage system transaction request entries were accumulated in the log.

1 40. (Previously Presented) The method of claim 26, further comprising:
2 modifying at least some messages in the swarm of messages based on the order in
3 which transaction request entries were accumulated in the log.

1 41. (Previously Presented) A file system, comprising:
2 a backup memory storing a plurality of file system transaction entries;
3 a first process that establishes a swarm of messages with respect to the file system
4 transaction entries and delivers the swarm of messages to the file system;
5 a second process that performs a parallel LOAD phase for a plurality of messages
6 in the swarm of messages where the LOAD phase is processed by commingling one or
7 more steps of the LOAD phase applied to each message of the swarm of messages; and
8 a third process that performs a MODIFY phase for at least some messages in the
9 swarm of messages, the MODIFY phase operating on messages based on the order in
10 which file system transaction entries were stored in the backup memory.

- 1 42. (Previously Presented) The file system of claim 41, further comprising:
2 a fourth process that determines whether a file system transaction entry corre-
3 sponds to a file system transaction that can be performed right away.
- 1 43. (Previously Presented) The file system of claim 42, wherein the fourth process,
2 in response to determining that the file system transaction can not be performed right
3 away, associates the file system transaction entry with a LOAD RETRY state until a prior
4 prerequisite transaction is performed.
- 1 44. (Previously Presented) A method, comprising:
2 storing a plurality of file system transaction entries in a backup memory;
3 establishing a swarm of messages with respect to the file system transac-
4 tion entries;
5 delivering the swarm of messages to a file system;
6 performing a parallel LOAD phase for a plurality of messages in the
7 swarm of messages where the LOAD phase is processed by commingling one or
8 more steps of the LOAD phase applied to each message of the swarm of mes-
9 sages; and
10 performing a MODIFY phase for at least some messages in the swarm of
11 messages, the MODIFY phase operating on messages based on the order in which
12 file system transaction entries were stored in the backup memory.
- 1 45. (Cancelled)
- 1 46. (Previously Presented) A method for operating a storage system, comprising:
2 storing a plurality of file system transaction entries in a backup memory;
3 establishing a swarm of messages with respect to the plurality of file system trans-
4 action entries;
5 delivering the swarm of messages to a file system; and

6 performing a parallel retrieval process by a processor within the storage system,
7 for a plurality of messages in the swarm of messages by processing the messages in an
8 arbitrary order, where the processor executes the messages in the arbitrary order by
9 commingling the processing of messages and steps of the retrieval process.

1 47. (Previously Presented) An operating system configured with a storage system, com-
2 prising:

3 a backup memory storing a plurality of transaction entries;

4 a first process that establishes a swarm of messages with respect to the transaction
5 entries and delivers the swarm of messages to the operating system;

6 a second process that performs a parallel LOAD phase for a plurality of messages
7 in the swarm of messages where the LOAD phase is processed by commingling one or
8 more steps of the LOAD phase applied to each message of the swarm of messages; and

9 a third process that performs a MODIFY phase for at least some messages in the
10 swarm of messages, the MODIFY phase operating on messages based on the order in
11 which the transaction entries were stored in the backup memory.

1 48. (Previously Presented) A system for replay of a backup memory in a storage system
2 having an operating system for managing transfer of data to and from an attached disk
3 array, the system comprising:

4 a log in the backup memory containing the storage system transaction entries ac-
5 cumulated after a consistency point at which time results of the storage system transac-
6 tion entries are committed to the disk array;

7 an initiator process that establishes a plurality of messages with respect to the
8 storage system transaction entries and delivers the plurality of messages to the operating
9 system; and

10 a parallel disk information-retrieval process in the operating system that is carried
11 out on the plurality of messages in parallel by the processor commingling steps of each
12 message of the plurality of messages.

1 49. (Previously Presented) A system for replay of a backup memory in a storage system
2 having an operating system for managing transfer of data to and from an attached disk
3 array, the system comprising:

4 a log in the backup memory containing the storage system transaction entries ac-
5 cumulated after a consistency point at which time results of the storage system transac-
6 tion entries are committed to the disk array;

7 an initiator process that establishes a plurality of messages with respect to the
8 storage system transaction entries and delivers the plurality of messages to the operating
9 system; and

10 a parallel disk information-retrieval process in the operating system that is carried
11 out in an arbitrary order to load the plurality of messages in parallel, where a processor
12 commingles steps of each message of the plurality of messages.